

DEEP ATTRACTOR NETWORK FOR SINGLE-MICROPHONE SPEAKER SEPARATION

Zhuo Chen Yi Luo Nima Mesgarani

Department of Electrical Engineering, Columbia University

ABSTRACT

Despite the overwhelming success of deep learning in various speech processing tasks, the problem of separating simultaneous speakers in a mixture remains challenging. Two major difficulties in such systems are the arbitrary source permutation and unknown number of sources in the mixture. We propose a novel deep learning framework for single channel speech separation by creating attractor points in high dimensional embedding space of the acoustic signals which pull together the time-frequency bins corresponding to each source. Attractor points in this study are created by finding the centroids of the sources in the embedding space, which are subsequently used to determine the similarity of each bin in the mixture to each source. The network is then trained to minimize the reconstruction error of each source by optimizing the embeddings. The proposed model is different from prior works in that it implements an end-to-end training, and it does not depend on the number of sources in the mixture. Two strategies are explored in the test time, K-means and fixed attractor points, where the latter requires no post-processing and can be implemented in real-time. We evaluated our system on Wall Street Journal dataset and show 5.49% improvement over the previous state-of-the-art methods.

Index Terms— Source separation, multi-talker, deep clustering, attractor network

1. INTRODUCTION

Despite the recent advances in deep learning methods for various speech processing tasks such as automatic recognition[1, 2, 3], enhancement[4, 5, 6], speech separation remains unresolved. Two main difficulties hinder the efficacy of deep learning algorithms to tackle the famous “cocktail party problem”[7].

The first difficulty is referred as the “permutation problem” [8]. Most neural network methods are trained to map the input signal to a unique target output which can be a label, a sequence, or regression reference. Permutation problem in speech separation arises due to the fact that the order of targets in the mixture is irrelevant. For example, in separating speakers A and B, both (A,B) and (B,A) solutions are acceptable. However, training the neural network with more than one target label per sample will produce conflicting gra-

dients and thus lead to convergence problem, because the order of targets cannot be determined beforehand. For instance, assigning speaker A to the first target position in (A,B) and (A,C) will cause confusions when the mixture is (B,C), since they both need to be in the second position for consistency.

The second problem in using neural network framework for speech separation is the output dimension mismatch problem. Since the number of sources in the mixture can vary, a neural network with fixed number of output nodes does not have the flexibility to separate arbitrary number of sources.

Two deep learning based methods have been proposed to resolve these problems, which are known as “deep clustering (DC)[8]” and “permutation invariant training (PIT)[9]”. In deep clustering, a network is trained to generate discriminative embedding for each time-frequency (T-F) bin with points belonging to the same source forced to be closer to each other. DC is able to solve both permutation and output dimension problem to produce the state of the art separation performance. The main drawback of DC is its inefficiency to perform end-to-end mapping, because the objective function is the affinity between the sources in the embedded space and not the separated signals themselves. Minimizing the separation error is done with an unfolding clustering system and a second network, which is trained iteratively and stage by stage to ensure convergence [10]. The PIT algorithm solves the permutation problem by pooling over all possible permutations for N mixing sources ($N!$ permutations), and use the permutation with lowest error to update the network. PIT was first proposed in [8], and was later shown to have comparable performance as DC [9]. However, PIT approach suffers the output dimension mismatch problem because it assumes a fixed number of sources. PIT also suffers from its computation efficiency, where the prediction window has to be much shorter than context window due to the inconsistency of the permutation both across and within sample segments.

In this work, we propose a novel deep learning framework which we refer to as the attractor network to solve the source separation problem. The term attractor refers to the well-studied perceptual effects in human speech perception which suggest that the brain circuits create perceptual attractors (magnets) that warp the stimulus space such that it draws the sound that is closest to it, a phenomenon that is called Perceptual Magnet Effect [11]. Our proposed model works on the same principle by forming a reference point (attractor)

for each source in the embedding space which draws all the T-F bins toward itself. Using the similarity between the embedded points and each attractor, a mask is estimated for each sources in the mixture. Since the mask is directly related to the attractor point, the proposed framework can potentially be extended to arbitrary number of sources without the permutation problem. Moreover, the mask learning enables a very efficient end-to-end training scheme and highly reduces the computation complexity compared with DC and PIT.

In Section 2, the proposed model is explained and discussed in more detail. In Section 3, we evaluate the performance of proposed system, and the conclusion is drawn in Section 4.

2. ATTRACTOR NEURAL NETWORK

2.1. Model

The neural network is trained to map the mixture sound X to a k dimensional embedding space, such that it minimizes the following objective function:

$$\mathcal{L} = \sum_{f,t,c} \|S_{f,t,c} - X_{f,t} \times M_{f,t,c}\|_2^2 \quad (1)$$

where S is the clean spectrogram (frequency $F \times$ time T) of C sources, X is the mixture spectrogram (frequency $F \times$ time T), and M is the mask formed to extract each source. The mask is estimated in the K dimensional embedding space of each T-F bin, represented by $V \in \mathbb{R}^{F \times T \times K}$:

$$M_{f,t,c} = \frac{1}{1 + \exp(\sum_k A_{c,k} \times V_{f,t,k})} \quad (2)$$

where $A \in \mathbb{R}^{C \times K}$ are the attractors for the C sources in the embedding space, learnt during training, which are defined as

$$A_{c,k} = \frac{\sum_{f,t} V_{k,f,t} \times Y_{c,f,t}}{\sum_{f,t} Y_{c,f,t}} \quad (3)$$

which $Y \in \mathbb{R}^{F \times T \times K}$ is the source membership function for each T-F bin, i.e., $Y_{t,f,c} = 1$ if source c has the highest energy at time t and frequency f compare to the other sources.

The objective function in Equation 1 consists of three parts. During training, we first compute an embedding V through a forward pass of the neural network for each given mixture. Then an attractor vector is estimated for each source using Equation 3. This can be done in several ways which we will elaborate in Section 2.2. The most straightforward method for attractor generation is to find the source centroid, as defined in Equation 3.

Next, we estimate a reconstruction mask for each source by finding the similarity of each T-F bin in the embedding space to each of the attractor vectors A , where the similarity metric is defined in Equation 2. This particular metric uses the inner product followed by a sigmoid function which

monotonically scales the masks between $[0, 1]$. Intuitively, if an embedding of a T-F bin is closer to one attractor means that it belongs to that source, and the resulting mask for that source will produce larger values for that T-F bin.

Finally, a standard $L2$ reconstruction error is used to generate the gradient, as shown in Equation 1. Therefore, the error for each source reflects the difference between the masked signal and the clean reference, forcing the network to optimize the global reconstruction error for better separation. We refer the proposed net as deep attractor network (DANet).

In comparison with previous methods, DANet network has two distinct advantages. Firstly, DANet removes the step-wise pre-training required in DC method to enable end-to-end training. Another big advantage of DANet arises from the flexibility in source dependent training, where the source-dependent knowledge could be easily incorporated by the attractor (e.g. speaker identity).

2.2. Estimation of attractor points

Attractor points can be estimated using various methods other than the average used in Equation 3. One possibility is to use weighted average. Since the attractors represents the source center of gravity, we can include only the embeddings of the most salient T-F bins, which leads to more robust estimation. We investigate this strategy by using an amplitude threshold in the estimation of the attractor. Alternatively, a neural network model may also be used to pick the representative embedding for each source, an idea which shares similarities with encoder-decoder attention networks [12, 13].

During test time, because the true assignment Y is unknown, we incorporate two strategies to form the attractor points. The first is similar to the strategy used in DC, where the centers are found using post K-means algorithm. The second method is based on the observation that the location of the attractors in the embedding space is relatively stable. This observation is shown in Figure 1, where each pair of dots corresponds to the attractor found for the two speakers in a given mixture. Figure 1 shows two principle stable attractor pairs for all the mixtures used, however, this observation needs to be tested in more depth and different tasks and datasets.

2.3. Relation with DC and PIT

The objective function of DC is shown in Equation 4, where Y is the indicator function which is equivalent to a binary mask, and V is the learnt embedding:

$$\mathcal{L} = \sum_{f,t} \left\| \sum_c Y Y^T - V V^T \right\|_2^2 \quad (4)$$

Since Y is orthonormal and constant for each mixture, by multiplying Y^T to both term, we can get an objective function that is a special case of the attractor network, as in Equation 5:

$$\mathcal{L} = \sum_{f,t} \left\| \sum_c Y^T - Y^T V V^T \right\|_2^2 \quad (5)$$

In Equation 5, $Y^T V$ can be viewed as an averaging step, where the embeddings are summed according to the label, and the resulted center is multiplied with the embedding matrix V to measure the similarity between each embedding and the center, and compared with the ground truth binary mask. Therefore, the transformed DC object can be viewed as a special example using binary mask as target and averaging-based attractor vector.

On the other hand, when the attractor vectors are considered as free parameters in the network, DANet reduces to a classification network [4, 5], and Equation 1 becomes a fully-connected layer. In this case, PIT becomes necessary since the mask has no information about the source and the problem of fixed output dimension arises. In contrast, the freedom of the network to form attractor points during the training allows the system to use the affinity between samples where no constraint is on the number of patterns found, therefore allowing the network to be independent of the number of sources. The flexibility of the network in choosing the attractor points is helpful even in two-source separation problem, because the two sources may have very different structures. As can be seen in Figure 2, our proposed method trained in speaker separation tasks has ended up finding 2 attractor pairs (4 points in the embedding space), which can be expected to increase in harder problems.

3. EVALUATION

3.1. Experimental setup

We evaluate our proposed model on the task of single-channel overlapped two-speaker separation. We use the corpus introduced in [8], which contains a 30 h training set and a 10 h validation set generated by randomly selecting utterances from different speakers in the Wall Street Journal (WSJ0) training set `si_tr_s`, and mixing them at various signal-to-noise ratios (SNR) randomly chosen between 0 dB and 10 dB. 5 h evaluation set is generated similarly as above, using utterances from 16 unseen speakers from `si_dt_05` and `si_et_05` in WSJ0 dataset. All data are resampled to 8 kHz to reduce computational and memory costs. The log spectral magnitude is served as input feature, computed using short-time Fourier transform (STFT) with 32 ms window length, 8 ms hop size, and the square root of hanning window.

The network contains 4 Bi-directional LSTM [14] layers with 600 hidden units in each layer. The embedding dimension is set to 20, resulting in a fully-connected feed-forward layer of 2580 hidden units (20×129) after the BLSTM layers. We split the input features into non-overlapping chunks of 100-frame length as the input to the network. RMSprop algorithm [15] is used for training with an exponential learning

rate decaying strategy, where the learning rate starts at 10^{-4} and ends at 3×10^{-6} . The total number of epochs was set to be 150, and we used the cost function in Equation 1 on the validation set for early stopping. The criteria for early stopping is no decrease in the loss function on validation set for 10 epochs. We constructed a Deep Clustering (DC) network with the same configuration which is used as the baseline.

We report the results in terms of signal-to-distortion ratio (SDR, which we define as scale-invariant SNR here), signal-to-artifacts ratio (SAR), and signal-to-interference ratio (SIR). The results are shown in Table 3.3

3.2. Separation examples

Figure 1 shows an example of mixture, the difference between the two speakers in the mixture, and the separated spectrograms of the two speakers using DANet. Also visualized in Figure 1 is the embeddings of the mixture projected onto its first Principal Components. In Figure 1, each point represents one T-F bin in the embedding space. Red and blue dots correspond to the T-F bins where speaker one or two have greater energy accordingly. The location of attractors are marked by x . It shows that two symmetric attractor centers are formed, each corresponding to one of the speakers in the mixture. A clear boundary can be observed in the figure, showing that the network successfully pulled the two speakers apart toward their corresponding attractor points.

Figure 2 shows the location of attractors for 10,000 mixture examples, mapped onto the 3-dimensional space for visualization purpose using Principal Component Analysis. It suggests that the network may have learned two attractor pairs (4 symmetric centers), marked by A1 and A2. This observation confirms our intuition of the DANet mentioned in Section , that DANet has the ability to discover different number of attractors in an unsupervised way, and therefore, form complex separation strategies. Although the task considered in this study is already challenging, one can imagine much more difficult separation scenarios, where the number of speakers in the mixture is large and can change over time. The ability of DANet to form new attractors may prove to be crucial in such cases, because any effort in pre-setting the number of mixing patterns, as done in methods such as PIT, will hinder the generalization ability of the network. Figure 2 also suggests that hierarchical clustering methods can be more suitable, where attractors can drive a hierarchical grouping of sources, allowing a better representation of audio signals. We will explore these issues in future work.

3.3. Results

Table 3.3 shows the evaluation results for different networks (example sounds can be found here [16]). Although the plain DANet already outperforms the DC baseline, adding a simple threshold on T-F samples included in the formation of the attractor yields further improved performance, presumably due

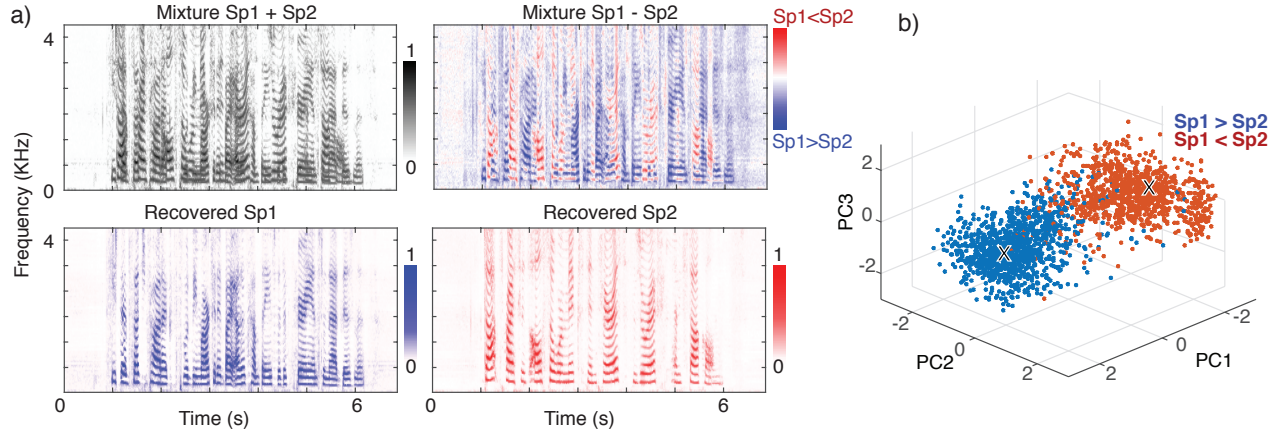


Fig. 1. Location of T-F bins in the embedded space. Each dot visualizes the first three principle components of one T-F bin, where colors distinguish the relative power of speakers, and the location of attractors is marked with X.




Fig. 2. Location of attractor points in the embedding space. Each dot corresponds to one of the 10000 mixtures sounds, visualized using the first three principal components. Two distinct attractor pairs are visible (denoted by A1 and A2).

to the increased influence of salient segments. On the other hand, the performance boost suggests that better attractor formation procedures can be utilized to estimate the attractors, such as joint optimizing of the network parameters.

In the last experiment in Table 3.3, a fixed pair of attention vector collected from the training data is used, corresponding to the A1 pair in Figure 2. This pre-set attractor is able to generalize well to the unseen mixtures and produced high quality separation, however it was slightly worse than the best model. Compared with K-means, this has the advantage that it can be easily implemented in real-time using a frame-by-frame pipeline. Based on this observation, when more attractors are required (e.g. in more complex tasks), a collection of attractor codebook and a simple classifier could be implemented for real-time processing.

| | GNSDR | GSAR | GSIR |
|------------------------|------------|-------------|-------------|
| DC | 9.1 | 9.5 | 22.2 |
| DANet | 9.4 | 10.1 | 18.8 |
| DANet-50% | 9.4 | 10.4 | 17.3 |
| DANet-70% | 9.6 | 10.3 | 18.7 |
| DANet-90% | 9.6 | 10.4 | 18.1 |
| DANet-90% [‡] | 9.5 | 10.4 | 17.8 |

Table 1. Evaluation metrics for networks with different configurations. The percentage suffixes stand for the salient weight threshold used during training. [‡]: Use fixed attractor points calculated by training set.

4. CONCLUSION

In this work, we proposed a novel neural network framework called deep attractor network for general source separation problem. The network forms attractor points in a high-dimensional embedding space of the signal, and the similarity between attractors and time-frequency embeddings are then converted into a soft separation mask. We showed that the proposed framework can perform end-to-end, real-time separation, is number-of-sources independent, and is more general comparing to deep clustering and classification based approaches. The experiment on two speaker separation task confirmed its efficacy and potential for extension to general source separation problems.

5. ACKNOWLEDGEMENT

The authors would like to thank Drs. John Hershey and Jonathan Le Roux of Mitsubishi Electric Research Lab for constructive discussions. This work was funded by a grant from National Institute of Health, NIDCD, DC014279, National Science Foundation CAREER Award, and the Pew Charitable Trusts.

6. REFERENCES

- [1] G.E. Dahl, D. Yu, L. Deng, and A. Acero, "Context-dependent pre-trained deep neural networks for large-vocabulary speech recognition," in *IEEE Trans. on Audio, Speech and Language Processing*, 2012, vol. 20.
- [2] S. Hašim, Senior A., and F. Beaufays, "Long short-term memory recurrent neural network architectures for large scale acoustic modeling," in *Interspeech 2014*. ISCA, 2014.
- [3] A. Senior, S. Hašim, F. C. Quitry, T. N. Sainath, and K. Rao, "Acoustic modelling with cd-ctc-smbr lstm rnns," in *ASRU 2015*. IEEE, 2015.
- [4] Zhuo Chen, Shinji Watanabe, Hakan Erdoğan, and John R Hershey, "Speech enhancement and recognition using multi-task learning of long short-term memory recurrent neural networks," ISCA, 2015.
- [5] Takaaki Hori, Zhuo Chen, Hakan Erdogan, John R Hershey, Jonathan Le Roux, Vikramjit Mitra, and Shinji Watanabe, "The merl/sri system for the 3rd chime challenge using beamforming, robust feature extraction, and advanced speech recognition," in *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*. IEEE, 2015, pp. 475–481.
- [6] X. Lu, Y. Tsao, S. Matsuda, and C. Hori, "Speech enhancement based on deep denoising autoencoder," 2013, pp. 436–440.
- [7] E Colin Cherry, "Some experiments on the recognition of speech, with one and with two ears," *The Journal of the acoustical society of America*, vol. 25, no. 5, pp. 975–979, 1953.
- [8] John R Hershey, Zhuo Chen, Jonathan Le Roux, and Shinji Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 31–35.
- [9] Dong Yu, Morten Kolbæk, Zheng-Hua Tan, and Jesper Jensen, "Permutation invariant training of deep models for speaker-independent multi-talker speech separation," *arXiv preprint arXiv:1607.00325*, 2016.
- [10] Yusuf Isik, Jonathan Le Roux, Zhuo Chen, Shinji Watanabe, and John R Hershey, "Single-channel multi-speaker separation using deep clustering," *arXiv preprint arXiv:1607.02173*, 2016.
- [11] Patricia K. Kuhl, "Human adults and human infants show a perceptual magnet effect," *Perception & psychophysics*, 50.2 (1991): 93-107.
- [12] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio, "Neural machine translation by jointly learning to align and translate," *arXiv preprint arXiv:1409.0473*, 2014.
- [13] Kyunghyun Cho, Aaron Courville, and Yoshua Bengio, "Describing multimedia content using attention-based encoder-decoder networks," *IEEE Transactions on Multimedia*, vol. 17, no. 11, pp. 1875–1886, 2015.
- [14] Sepp Hochreiter and Jürgen Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [15] Tijmen Tieleman and Geoffrey Hinton, "Lecture 6.5-rmsprop," *COURSERA: Neural networks for machine learning*, 2012.
- [16] "naplab.ee.columbia.edu/anet," .